

Stanford Community Advisory Committee:

SHORTLIFFE, Edward H., M.D., Ph.D. (Chairman)
Principal Investigator - SUMEX
Medical School Office Building, Rm. X271
Stanford University Medical Center
Stanford, California 94305
(415) 723-6979

FEIGENBAUM, Edward A., Ph.D.
Heuristic Programming Project
Department of Computer Science
Margaret Jacks Hall
Stanford University
Stanford, California 94305
(415) 723-4879

LEVINTHAL, Elliott C., Ph.D.
Departments of Mechanical and Electrical Engineering
Building 530
Stanford University
Stanford, California 94305
(415) 723-9037

Appendix E

Scientific Subproject Abstracts

The following are brief abstracts of our collaborative research projects.

National AIM Project: ATTENDING Project:
 A Critiquing Approach to
 Expert Computer Advice

Principal Investigator: Perry L. Miller, M.D., Ph.D.
 Department of Anesthesiology
 Yale University School of Medicine
 New Haven, CT 06510
 (203) 785-2802

Our project is exploring the "critiquing" approach to bringing computer-based advice to the practicing physician.

Critiquing is a different approach to the design of artificial intelligence based expert systems. Most medical expert systems attempt to simulate a physician's decision-making process. As a result, they have the clinical effect of trying to tell a physician what to do: how to practice medicine. In contrast, a critiquing system first asks the physician how he contemplates approaching his patient's care, and then critiques that plan. In the critique, the system discusses any risks or benefits of the proposed approach, and of any other approaches which might be preferred. It is anticipated that the critiquing approach may be particularly well suited for domains, like medicine, where decisions involve a great deal of *subjective* judgment.

To date, several prototype critiquing systems have been developed in different medical domains:

1. ATTENDING, the first system to implement the critiquing approach, critiques anesthetic management.
2. HT-ATTENDING critiques the pharmacologic management of essential hypertension.
3. VQ-ATTENDING critiques aspects of ventilator management.
4. PHEO-ATTENDING critiques the laboratory and radiologic workup of a patient for a suspected pheochromocytoma.

In addition, a domain-independent system, ESSENTIAL-ATTENDING, has been developed to facilitate the implementation of critiquing systems in other domains.

PUBLICATIONS

1. Miller, P.L. (Ed.): Selected Topics in Medical Artificial Intelligence. New York: Springer-Verlag (in press).
2. Rennels, G.D., Miller, P.L.: Artificial intelligence research in anesthesia and intensive care. Journal of Clinical Monitoring (in press).
3. Miller, P.L., Rennels, G.D.: Prose generation from expert systems: An applied computational linguistics approach. AI Magazine (in press).
4. Rennels, G.D., Shortliffe, E.H., Stockdale, F.E., Miller, P.L.: A computational model of reasoning from the clinical literature. Computer Methods and Programs in Biomedicine 24:139-149, 1987.
5. Rennels, G.D., Shortliffe, E.H., Stockdale, F.E., Miller, P.L.: A structured representation of the clinical literature and its use in a medical management advice system. Bulletin du Cancer 74:215-220, 1987.

6. Miller, P.L., Barwick, K.W., Morrow, J.S., Powsner, S.M., Riely, C.A.: Semantic relationships and medical bibliographic retrieval: A preliminary assessment. *Computers and Biomedical Research* 21:64-77, 1988.
7. Miller, P.L., Morrow, J.S., Powsner, S.M., Riely, C.A.: Semantically assisted medical bibliographic retrieval: An experimental computer system. *Bulletin of the Medical Library Association* 76:131-136, 1988.
8. Miller, P.L.: Exploring the critiquing approach: Clinical practice-based feedback by computer. *Biomedical Measurement, Informatics and Control* (in press).
9. Rennels, G.D., Shortliffe, E.H., Stockdale, F.E., Miller, P.L.: A computational model of reasoning from the clinical literature. *The AI Magazine* (accepted pending revision).
10. Miller, P.L., Fisher, P.R.: Causal models in medical artificial intelligence. *Proceedings of the Eleventh Symposium on Computer Applications in Medical Care, Washington, D.C., November 1987*, pp. 17-22.
11. Powsner, S.M., Barwick, K.W., Morrow, J.S., Riely, C.A., Miller, P.L.: Coding semantic relationships for medical bibliographic retrieval: A preliminary study. *Proceedings of the Eleventh Symposium on Computer Applications in Medical Care, Washington, D.C., November 1987*, pp. 108-112.

Stanford Project: BBICU -- BLACKBOARD APPLICATIONS
IN THE INTENSIVE CARE UNIT

Principal Investigator: Adam Seiver, M.D.
Department of Surgery
Palo Alto Veterans Medical Center
Palo Alto, Calif

Lawrence Fagan, M.D., Ph.D.
Department of Medicine
Stanford University

Barbara Hayes-Roth, Ph.D.
Department of Computer Science
Stanford University

We are designing a data-interpretation and therapy-planning system for the intensive care unit (ICU). Fundamental research issues in temporal reasoning are associated with the ICU application area including assimilation of incoming data, representation of time-oriented intervals, and description of ongoing physiological processes [Fagan 84]. In addition, in ICUs of the 1990s, many more physiological measurements will need to be collected at frequent intervals, and increased access to the current medical record in coded format will be possible. Processing of incoming data will have to be opportunistic, selecting from a number of models that have different computational requirements and accuracy. We will use a blackboard architecture, known as BB1, that has evolved from earlier work on protein-structure elucidation and construction layout [Hayes-Roth 85]. BB1 is particularly well suited for the ICU project because it maintains separate blackboards for domain and control knowledge.

Although we have selected the blackboard structure as the organizing principle, many knowledge representation issues remain. First, we must represent the structure and function (anatomy and physiology) of the respiratory system. By characterizing the pathophysiology in terms of generic faults, we will create a more flexible means to diagnose problems in unusual situations -- in contrast to the phenomenological rules used in earlier systems.

Second, we must coalesce quantitative and qualitative models. The physiology of the respiratory and cardiac systems have been modeled in detail, but it is impractical to base the entire reasoning process on complex mathematical equations. Instead, we are developing methods to transform quantitative models into simpler formulations. We must make explicit the simplifying assumptions and associate them with their corresponding models, so that we can select an applicable model for situations of varying complexity.

Using the approach of our project for planning treatments for cancer patients [Langlotz 87], we will use strategic knowledge to create patient-specific specializations of standard treatment plans. We will use decision analytic methods to evaluate and explain the various treatment options available at any point in time. The long-term goal of this project is to embed the decision-making components within the data management tasks of the ICU.

REFERENCES

1. Fagan, L., Kunz, J., Feigenbaum, E, and Osborn, J. Adapting a rule-based system for a monitoring task, in *Rule Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project*, B. Buchanan and E. Shortliffe (eds.). Reading, MA: Addison-Wesley Publishing Co., 1984.

2. Hayes-Roth, B. A Blackboard architecture for control. (Artificial Intelligence) 26:251-321, 1985.
3. Langlotz, C., Fagan, L., Tu, S., Sikic, B., and Shortliffe, E. A therapy planning architecture that combines decision theory and artificial intelligence techniques. *Computers and Biomedical Research* 20:279-303, 1987.

National AIM Project: DECISION SUPPORT FOR
 TIME-VARYING CLINICAL PROBLEMS

Principal Investigator: Lawrence Widman, M.D., Ph.D.
 Division of Cardiology
 University of Texas Health Science Center
 San Antonio, Texas 78284

Time-varying systems, which include many areas of medicine, science, economics, and business, can be described mathematically by differential equations. They are distinct from the pattern-matching and logic-based domains dealt with so successfully by existing expert system methods, because they can include feedback relationships. It is generally felt that they are best approached by enhancement of existing methods for deep model-based reasoning.

The goal of this project is to develop AI methods for capturing and using knowledge about time-varying systems. The strategy is to address general problems in model-based knowledge representation and reasoning. The intermediate objective is to develop methods which are powerful enough to work in selected realistic situations yet are general enough to be transportable to other, unrelated knowledge domains.

The tactical approach is to work on well-defined yet complex and interesting problems in the medical domain. We have, therefore, selected the human cardiovascular system as our prototype of a time-varying system, and are developing methods for representing and reasoning about its mechanical and electrical activities in the normal and diseased states.

REFERENCES

1. Widman, L.E. Knowledge-Based Fault Identification and "What If" Simulation in Symbolic Dynamic Systems Models. Proceedings of the 1988 Society for Computer Simulation Multiconference on AI and Simulation, pp. 89-94, 1988.
2. Widman, L.E., Lee, Y.-B., and Y.-H. Pao. Diagnosis of Causal Medical Models by Semi-Quantitative Reasoning. In: Miller, P.L. (ed.). Topics in Medical Artificial Intelligence, Springer-Verlag (in press)
3. Widman, L.E. Semi-Quantitative "Close Enough" Systems Dynamics Models: An Alternative to Qualitative Simulation. In: Widman, L.E., Loparo, K.A., and N.R. Nielsen (eds.). Artificial Intelligence, Simulation, and Modeling. John Wiley & Sons, New York (in press).
4. Widman, L.E., Loparo, K.A., and N.R. Nielsen (eds.). Artificial Intelligence, Simulation, and Modeling. John Wiley & Sons, New York (in press).
5. Widman, L.E. and G.L. Freeman. A-to-D Conversion from Paper Records with a Desktop Scanner and a Microcomputer. (submitted to the American Journal of Physiology).

Stanford Project: GUIDON/NEOMYCIN --
KNOWLEDGE ENGINEERING
FOR TEACHING MEDICAL DIAGNOSIS

Principal Investigators: William J. Clancey, Ph.D.
701 Welch Road
Department of Computer Science
Stanford University
Palo Alto, California 94304
(415) 723-1997 (CLANCEY@SUMEX-AIM)

Bruce G. Buchanan, Ph.D.
Computer Science Department
Stanford University

The GUIDON/NEOMYCIN Project is a research program devoted to the development of a knowledge-based tutoring system for application to medicine. The key issue for the GUIDON/NEOMYCIN project is to develop a program that can provide advice similar in quality to that given by human experts, modeling how they structure their knowledge as well as their problem-solving procedures. The consultation program using this knowledge is called NEOMYCIN. NEOMYCIN's knowledge base, designed for use in a teaching application, is the subject material used by a family of instructional programs referred to collectively as GUIDON2. The problem-solving procedures are developed by running test cases through NEOMYCIN and comparing them to expert behavior. Also, we use NEOMYCIN as a test bed for the explanation capabilities incorporated in our instructional programs.

The goal of current work is to construct a knowledge-based tutoring system that teaches diagnostic strategies explicitly. By strategy, we mean plans for establishing a set of possible diagnoses, focusing on and confirming individual diagnoses, gathering data, and processing new data. The tutorial program has capabilities to recognize these plans, as well as to articulate strategies in explanations about how to do diagnosis. The strategies represented in the program, modeling techniques, and explanation techniques are wholly separate from the knowledge base, so that they can be used with many medical (and non-medical) domains. That is, the target program will be able to be tested with other knowledge bases, using system-building tools that we provide.

SOFTWARE AVAILABLE ON SUMEX

GUIDON--A system developed for intelligent computer-aided instruction. Although it was developed in the context of MYCIN's infectious disease knowledge base, the tutorial rules will operate upon any EMYCIN knowledge base.

NEOMYCIN--A consultation system derived from MYCIN, with the knowledge base greatly extended and reconfigured for use in teaching. In contrast with MYCIN, diagnostic procedures, common sense facts, and disease hierarchies are factored out of the basic finding/disease associations. The diagnostic procedures are abstract (not specific to any problem domain) and model human reasoning, unlike the exhaustive, top-down approach implicit in MYCIN's medical rules. This knowledge base is used in the GUIDON2 family of instructional programs, being developed on D-machines.

REFERENCES

1. Clancey, W.J.: **Knowledge-Based Tutoring: The GUIDON Program**, Cambridge: The MIT Press, 1987.
2. Clancey, W.J.: *Methodology for building an intelligent tutoring system*. In Kintsch, Polson, and Miller, (Eds.), **METHODS AND TACTICS IN COGNITIVE SCIENCE**. L. Erlbaum Assoc., Hillsdale, NJ. 1984. (Also STAN-CS-81-894, HPP 81-18)
3. Clancey, W.J.: *Acquiring, representing, and evaluating a competence model of diagnosis*. In Chi, Glaser, and Farr (Eds.), **THE NATURE OF EXPERTISE**. 1988.

National AIM Project: INTERNIST-I/QMR (Quick Medical Reference)

Principal Investigators:

INTERNIST-QMR PROJECT:
Randolph A. Miller, M.D. (RMILLER@SUMEX-AIM)
Jack D. Myers, M.D. (MYERS@SUMEX-AIM)
Fred E. Masarie, Jr., M.D. (MASARIE@SUMEX-AIM)
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
(412) 648-3190

The major goal of the INTERNIST-I/QMR Project is to produce a reliable and adequately complete diagnostic consultative program in the field of internal medicine. Although this program is intended primarily to aid skilled internists in complicated medical problems, the program may have spin-offs as a diagnostic and triage aid to physicians' assistants, rural health clinics, military medicine and space travel. In the design of INTERNIST-I and QMR, we have attempted to model the creative, problem-formulation aspect of the clinical reasoning process, and to design program components that can help physicians at key points in this process. The program can compose differential diagnoses, dynamically, on the basis of clinical evidence. During the course of an INTERNIST-I consultation, it is not uncommon for a number of such conjectured problem foci to be proposed and investigated. With the QMR program, the physician is given more control over the direction of the problem-solving activities than was possible with INTERNIST-I. QMR is broader in scope than INTERNIST-I in that it provides quick and efficient access to the INTERNIST-I/QMR knowledge base to provide low and intermediate level informational support for physicians' decision-making, in addition to providing consultative advice.

SOFTWARE AVAILABLE ON SUMEX

Versions of the INTERNIST-I KB are available for experimental use, with permission from the developers. In the past year, copies have been shared with Drs. Gregory Cooper and Homer Chin of Dr. Shortliffe's group at Stanford. The INTERNIST-I and QMR knowledge bases are Copyright 1986-88 by the University of Pittsburgh. Quick Medical Reference and QMR are registered trademarks of the University of Pittsburgh. The project continues to be oriented primarily towards research and development; hence, a stable production version of the system is not yet available for general use. QMR has been shared on a restricted basis with a limited number of academic colleagues, who have agreed to give the QMR development team feedback on the program's strengths and weaknesses.

REFERENCES

1. Myers JD. The Background of INTERNIST-I and QMR. In: Proceedings of the History of Medical Informatics Conference. National Library of Medicine. pp. 195-197, November 1987.
2. Masarie, F.E., Miller, R.A. Medical Subject Headings and Medical Terminology: An analysis of terminology used in hospital charts. Bulletin of the Medical Library Association, 1987; 75:89-94.
3. Masarie FE, Miller RA. INTERNIST-I to Quick Medical Reference (QMR): Transition from a mainframe to a microcomputer. Proceedings Ninth Annual IEEE/Engineering in Biology and Biology Society. IEEE Press. pp. 1521- 1522, November 1987.
4. Parker, RC, Miller RA. Using causal knowledge to create simulated

- patient cases: The CPCS project as an extension of INTERNIST-I. Proceedings of the Eleventh Annual Symposium on Computer Applications in Medical Care. pp. 473-480, November 1987.
5. Bankowitz RA, Blumenfeld BH, Miller RA, et al. User variability in abstracting and entering printed case histories with Quick Medical Reference (QMR). Proceedings of the Eleventh Annual Symposium on Computer Applications in Medical Care. pp. 68-73, November 1987.
 6. Masarie FE, Miller RA. Quick Medical Reference (QMR): An information management tool for clinical diagnosis. IN: Critical Reviews in Medical Informatics. Boca Raton, Florida, CRC Press, 1988.
 7. Parker RC, Miller RA. Using causal knowledge to create simulated patient cases: The CPCS project as an extension of INTERNIST-I. IN: Miller PL (ed) Topics in Medical Artificial Intelligence. Computers and Medicine Series, Springer-Verlag, New York, 1988.
 8. Challinor SM, McNeil MA, Bankowitz RA, et al. Evaluation of a Computer-Assisted General Medicine Diagnostic Consultation Service. Abstract presented at the 11th Annual SGIM Meeting, Washington, D.C., April 27-29, 1988.
 9. Miller RA. From Automated Medical Records to Expert System Knowledge Bases: Common Problems in Representing and Processing Patient Data. Topics in Health Record Management. 75(3):23-36, 1987
 10. Miller RA. Computer-based Diagnostic Decisionmaking. Medical Care. 25(12):S148-S152, 1987.

National AIM Project: MENTOR -- MEDICAL EVALUATION OF
THERAPEUTIC ORDERS

Principal Investigators: Stuart Speedie, Ph.D. (SPEEDIE@SUMEX-AIM)
School of Pharmacy
University of Maryland
20 N. Pine Street
Baltimore, Maryland 21201
(301) 528-7650

Terrence F. Blaschke, M.D. (BLASCHKE@SUMEX-AIM)
Department of Medicine
Division of Clinical Pharmacology
Stanford University Medical Center
Stanford, California 94305

The goal of the MENTOR project is to implement and begin evaluation of a computer-based methodology for reducing therapeutic misadventures. The project uses an on-line expert system to continuously monitor the drug therapy of individual patients and generate specific warnings of potential and/or actual unintended effects of therapy. The appropriate patient information is automatically acquired through interfaces to a hospital information system. This data is monitored by a system that is capable of employing complex chains of reasoning to evaluate therapeutic decisions and arrive at valid conclusions in the context of all information available on the patient. The results reached by the system are fed back to the responsible physicians to assist future decision making.

Specific objectives of this project include:

1. Implement a prototype computer-based expert system to continuously monitor in-patient drug therapy that uses a modular medical knowledge base and a separate inference engine to apply the knowledge to specific situations.
2. Select a small number of important and frequently occurring drug therapy problems that can lead to therapeutic misadventures and construct a comprehensive knowledge base necessary to detect these situations.
3. Evaluate the prototype MENTOR system with respect to its impact on the on the physicians' therapeutic decision making as well as its effects on the patient in terms of specific mortality and morbidity measures.

The work in this project builds on the extensive previous work in drug monitoring done by these investigators in the Division of Clinical Pharmacology at Stanford and the University of Maryland School of Pharmacy.

Stanford Project: MOLGEN -- AN EXPERIMENT PLANNING SYSTEM
 FOR MOLECULAR GENETICS

Principal Investigators: Edward A. Feigenbaum, Ph.D.
 Department of Computer Science
 Stanford University

 Charles Yanofsky, Ph.D. (YANOFSKY@SUMEX-AIM)
 Department of Biology
 Stanford University
 Stanford, California 94305
 (415) 725-3815

The MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. Our current research concentrates on scientific discovery within the subdomain of regulatory genetics. We desire to explore the methodologies scientists use to modify, extend, and test theories of genetic regulation, and then emulate that process within a computational system.

Theory or model formation is a fundamental part of scientific research. Scientists both use and form such models dynamically. They are used to predict results (and therefore to suggest experiments to test the model) and also to explain experimental results. Models are extended and revised both as a result of logical conclusions from existing premises and as a result of new experimental evidence.

Theory formation is a difficult cognitive task, and one in which there is substantial scope for intelligent computational assistance. Our research is toward building a system which can form theories to explain experimental evidence, can interact with a scientist to help to suggest experiments to discriminate among competing hypotheses, and can then revise and extend the growing model based upon the results of the experiments.

The MOLGEN project has continuing computer science goals of exploring issues of knowledge representation, problem-solving, discovery, and planning within a real and complex domain. The project operates in a framework of collaboration between the Heuristic Programming Project (HPP) in the Computer Science Department and various domain experts in the departments of Biochemistry, Medicine, and Biology. It draws from the experience of several other projects in the HPP which deal with applications of artificial intelligence to medicine, organic chemistry, and engineering.

SOFTWARE AVAILABLE ON SUMEX

SPEX system for experiment design.

UNITS system for knowledge representation and acquisition.

SEQ system for nucleotide sequence analysis.

REFERENCES

1. Friedland, P.E.: *Knowledge-based experiment design in molecular genetics*, (Ph.D. thesis). Stanford Computer Science Report, STAN-CS-79-771.

2. Friedland, P.E. and Iwasaki, Y.: *The concept and implementation of skeletal plans*, Journal of Automated Reasoning, 1(2):161-208, 1985.
3. Friedland, P.E. and Kedes, L.: *Discovering the secrets of DNA*, Communications of the ACM, 28(11):1164-1186, November, 1985.
4. Stefik, M.J.: *An examination of a frame-structured representation system*, Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 845-852.
5. Stefik, M.J.: *Planning with constraints*, (Ph.D. thesis). Stanford Computer Science Report, STAN-CS-80-784, March, 1980.
6. Karp, P., and D. Wilkins: *An Analysis of the Deep/Shallow Distinction for Expert Systems*. Stanford University Knowledge Systems Laboratory Report KSL-86-32, 1986.
7. Karp, P., and P. Friedland: *Coordinating the Use of Qualitative and Quantitative Knowledge in Declarative Device Modeling*. Stanford University Knowledge Systems Laboratory Report KSL-87-09, 1987.
8. Round, A.: *QSOPS: A Workbench Environment for the Qualitative Simulation of Physical Processes*. Stanford University Knowledge Systems Laboratory Report KSL-87-37, 1987.
9. Karp, P.: *A Process-Oriented Model of Bacterial Gene Regulation*, Stanford University Knowledge Systems Laboratory Technical Report KSL-88-18.

Stanford Project: ONCOCIN -- KNOWLEDGE ENGINEERING FOR
ONCOLOGY CHEMOTHERAPY CONSULTATION

Principal Investigator: Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University Medical Center
Medical School Office Building
Stanford, California 94305
(415) 723-6979 (SHORTLIFFE@SUMEX-AIM)

Project Director: Dr. Lawrence M. Fagan (FAGAN@SUMEX-AIM)

The ONCOCIN Project is overseen by a collaborative group of physicians and computer scientists who are developing an intelligent system that uses the techniques of knowledge engineering to advise oncologists in the management of patients receiving cancer chemotherapy. The general research foci of the group members include knowledge acquisition, inexact reasoning, explanation, and the representation of time and of expert thinking patterns. Much of the work developed from research in the 1970's on the MYCIN and EMYCIN programs, early efforts that helped define the group's research directions for the coming decade. MYCIN and EMYCIN are still available on SUMEX for demonstration purposes.

The prototype ONCOCIN system is in limited experimental use by oncologists in the Stanford Oncology Clinic. Thus, much of the emphasis of this research has been on human engineering so that the physicians will accept the program as a useful adjunct to their patient care activities. ONCOCIN has generally been well-accepted since its introduction, and we are now testing a version of the program which runs on professional workstations (rather than the central SUMEX computer) so that it can be implemented and evaluated at sites away from the University.

Significant extensions to the basic ONCOCIN project have included the OPAL oncology knowledge acquisition system, the PROTEGE system for description of knowledge acquisition tools for clinical trial application areas, and the ONYX system for strategic therapy planning.

SOFTWARE AVAILABLE ON SUMEX

- MYCIN-- A consultation system designed to assist physicians with the selection of antimicrobial therapy for severe infections. It has achieved expert level performance in formal evaluations of its ability to select therapy for bacteremia and meningitis. Although MYCIN is no longer the subject of an active research program, the system continues to be available on SUMEX for demonstration purposes and as a testing environment for other research projects.
- EMYCIN-- The "essential MYCIN" system is a generalization of the MYCIN knowledge representation and control structure. It is designed to facilitate the development of new expert consultation systems for both clinical and non-medical domains.
- ONCOCIN-- This system is in clinical use but requires Lisp machines to be run. Much of the knowledge in the domain of cancer chemotherapy is already well-specified in protocol documents, but expert judgments also need to be understood and modeled.

REFERENCES

1. Shortliffe, E.H., Scott, A.C., Bischoff, M.B., Campbell, A.B., van Melle, W. and Jacobs, C.D.: *ONCOCIN: An expert system for oncology protocol management*. Proc. Seventh IJCAI, pp. 876-881, Vancouver, B.C., August, 1981.
2. Duda, R.O. and Shortliffe, E.H.: *Expert systems research*. Science 220:261-268, 1983.
3. Langlotz, C.P. and Shortliffe, E.H.: *Adapting a consultation system to critique user plans*. Int. J. Man-Machine Studies 19:479-496, 1983.
4. Bischoff, M.B., Shortliffe, E.H., Scott, A.C., Carlson, R.W. and Jacobs, C.D.: *Integration of a computer-based consultant into the clinical setting*. Proceedings 7th Annual Symposium on Computer Applications in Medical Care, pp. 149-152, Baltimore, Maryland, October 1983.
5. Hickam, D.H., Shortliffe, E.H., Bischoff, M.B., Scott, A.C., Jacobs, C.D.: A study of the treatment advice of a computer-based cancer chemotherapy protocol advisor (Memo KSL-85-21). *Annals of Internal Medicine* 103(6 pt 1):928-936 (1985).
6. Kent, D.L., Shortliffe, E.H., Carlson, R.W., Bischoff, M.B., Jacobs, C.D.: Improvements in data collection through physician use of a computer-based chemotherapy treatment consultant (Memo KSL-85-22). *Journal of Clinical Oncology* 3:1409-1417, 1985.
7. Tsuji, S. and Shortliffe, E.H.: Graphical access to a medical expert system: I. Design of a knowledge engineer's interface (Memo KSL-85-11). *Meth. Inf. Med.*, 25:62-70, 1986.
8. Lane, C.D., Differding, J.C., Shortliffe, E.H.: Graphical access to a medical expert system: II. Design of an interface for physicians (Memo KSL-85-15). *Meth. Inf. Med.*, 25:143-150, 1986.
9. Shortliffe, E.H. Medical expert systems: Knowledge tools for physicians. Memo KSL-86-52. Special issue on Medical Informatics, *West. J. Med.* 145:830-839, 1986.
10. Walton, J.D., Musen, M.A., Combs, D., Lane, C.D., Shortliffe, E.H., and Fagan, L.M. Graphical access to medical expert systems: III. Design of a knowledge-acquisition environment. Memo KSL-85-30. *Methods of Information in Medicine*, 26:78-88, 1987.
11. Langlotz, C.P., Fagan, L.M., Tu, S.W., Sikic, B.I., and Shortliffe, E.H. A therapy planning architecture that combines decision theory and artificial intelligence techniques. KSL-85-55. *Computers in Biomedical Research*, 20:279-303, 1987.
12. Musen, M.A., Fagan, L.M., Combs, D.M., and Shortliffe, E.H. Use of a domain model to drive an interactive knowledge-editing tool (Memo KSL-86-24). *International Journal of Man-Machine Studies* 26(1):105-121, 1987.

National AIM Project: Computer-Aided Diagnosis of
Lymph Node Pathology (PATHFINDER)

Principal Investigator: Bharat Nathwani, M.D.
Department of Pathology
HMR 204
2025 Zonal Avenue
University of Southern California
School of Medicine
Los Angeles, California 90033
(213) 226-7064 (NATHWANI@SUMEX-AIM)

The Pathfinder Project is centered on the construction of an expert system for assisting pathologists with the diagnosis of tissue pathology. Pathfinder research is focused on the domain of lymph node pathology. The project is based at the University of Southern California in collaboration with the Stanford University Medical Computer Science Group. Ongoing AIM research has been addressing fundamental problems of knowledge representation, reasoning strategies, user modeling, explanation, and user acceptance. The central focus of research has been on developing tractable knowledge acquisition and reasoning methods for probabilistic reasoning. In addition to the ongoing theoretical research, the team have implemented an expert system program and a tool for knowledge acquisition. The Pathfinder expert system provides diagnostic advice on eighty common benign and malignant diseases of the lymph nodes based on 150 histologic features. Currently a more sophisticated knowledge base that represents probabilistic dependencies is being developed. A probabilistic knowledge acquisition tool has been implemented on the Macintosh that enables a user to assess probabilistic relationships by graphically noting similarities and differences between diseases hypotheses.

SOFTWARE AVAILABLE ON SUMEX

Pathfinder -- A pilot version of the Pathfinder program is available for experimentation on the DEC 2060 computer. This version is an early version of the program that has not been completely tested.

REFERENCES

1. Horvitz, E.J., Heckerman, D.E., Nathwani, B.N. and Fagan, L.M.: *Diagnostic Strategies in the Hypothesis-directed Pathfinder System, Node Pathology*. HPP Memo 84-13. Proceedings of the First Conference on Artificial Intelligence Applications, Denver, Colorado, Dec., 1984.
2. Heckerman, D. E., and Horvitz, E. J., "The Myth of Modularity in Rule-based Systems," in *Uncertainty in Artificial Intelligence*, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland, New York, 1987.
3. Horvitz, E.J., Heckerman, D.E., Nathwani, B.N. and Fagan, L.M.: *The Use of a Heuristic Problem-solving Hierarchy to Facilitate the Explanation of Hypothesis-directed Reasoning*. KSL Memo 86-2. Proceedings of MedInfo, Washington D.C., October, 1986.
4. Horvitz, E. J., "Toward a Science of Expert Systems," Invited Paper, *Computer Science and Statistics: Proceedings of the 18th Symposium on the Interface*, American Statistical Association, March, 1986, pgs. 45-52.
5. Heckerman, D.E., "An Axiomatic Framework for Belief Updates," in

- Uncertainty in Artificial Intelligence, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland; New York, 1987.
6. * Heckerman, D. E., and Horvitz, E. J., "The Myth of Modularity in Rule-based Systems," in Uncertainty in Artificial Intelligence, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland, New York, 1987.
 7. * Heckerman, D.E., and Horvitz, E.J., "On the expressiveness of rule-based systems for reasoning under uncertainty," Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, July, 1987.
 8. * Horvitz, E. J., Heckerman, D. E., Langlotz, C. P., "A framework for comparing alternative formalisms for plausible reasoning," Proceedings of the AAAI," August, 1986, Morgan Kaufman, Los Altos, CA, 1986,
 9. * Horvitz, E.J., Breese, J.S., Henrion, M., Decision Theory in Expert Systems and Artificial Intelligence, International Journal of Approximate Reasoning, Elsevier, N.Y. July, 1988.
 10. * Heckerman, D.E., An Evaluation of Three Scoring Schemes, Proceedings of the 4th AAAI Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN., (to appear August 1988).
 11. * Horvitz, E.J., A Multiattribute Utility Approach to Inference Understandability and Explanation, Tech. Report, KSL-28-87, Knowledge Systems Laboratory, Stanford, California, March, 1987.

Stanford Project: PROTEAN Project

Principal Investigators: Oleg Jardetzky
 (JARDEZKY@SUMEX-AIM.STANFORD.EDU)
 Nuclear Magnetic Resonance Lab, School of Medicine
 Stanford University Medical Center
 Stanford, California 94305

 Bruce G. Buchanan, Ph.D.
 (BUCHANAN@SUMEX-AIM.STANFORD.EDU)
 Computer Science Department
 Stanford University
 Stanford, California 94305

The goals of this project are related both to biochemistry and artificial intelligence: (a) use existing AI methods to aid in the determination of the 3-dimensional structure of proteins in solution (not from x-ray crystallography proteins), and (b) use protein structure determination as a test problem for experiments with the AI problem-solving structure known as the Blackboard Model. Empirical data from nuclear magnetic resonance (NMR) and other sources may provide enough constraints on structural descriptions to allow protein chemists to bypass the laborious methods of crystallizing a protein and using X-ray crystallography to determine its structure. This problem exhibits considerable complexity, yet there is reason to believe that AI programs can be written that reason much as experts do to resolve these difficulties. A prototype knowledge-based system assembles major secondary structures of a protein into families of structures compatible with a given set of distance constraints under the control of an explicit assembly strategy. Structures can also be refined at the atomic level of detail using constraints within secondary structures and between amino acid side chains to further restrict the 3-dimensional structure found. By generalizing this approach to the assembly of arrangements of objects subject to constraints, we have developed a language for specifying actions and control for problem solving in similar problem domains.

REFERENCES

1. Altman, R. and Jardetzky, O.: *New strategies for the determination of macromolecular structures in solution*. Journal of Biochemistry (Tokyo), Vol. 100, No. 6, p. 1403-1423, 1986.
2. Altman, R. and Buchanan, B.G.: *Partial Compilation of Control Knowledge*. To appear in: Proceedings of the AAAI 1987.
3. Brinkley, J., Cornelius, C., Altman, R., Hayes-Roth, B. Lichtarge, O., Duncan, B., Buchanan, B.G., Jardetzky, O.: *Application of Constraint Satisfaction Techniques to the Determination of Protein Tertiary Structure*. Report KSL-86-28, Department of Computer Science, 1986.
4. Brinkley, James F., Buchanan, Bruce G., Altman, Russ B., Duncan, Bruce S., Cornelius, Craig W.: *A Heuristic Refinement Method for Spatial Constraint Satisfaction Problems*. Report KSL 87-05, Department of Computer Science.
5. Buchanan, B.G., Hayes-Roth, B., Lichtarge, O., Altman, A., Brinkley, J., Hewett, M., Cornelius, C., Duncan, B., Jardetzky, O.: *The Heuristic Refinement Method for Deriving Solution Structures of Proteins*. Report KSL-85-41. October 1985.

6. Garvey, Alan, Cornelius, Craig, and Hayes-Roth, Barbara: *Computational Costs versus Benefits of Control Reasoning*. Report KSL 87-11, Department of Computer Science.
7. Hayes-Roth, B.: *The Blackboard Architecture: A General Framework for Problem Solving?* Report HPP-83-30, Department of Computer Science, Stanford University, 1983.
8. Hayes-Roth, B.: *BB1: An Environment for Building Blackboard Systems that Control, Explain, and Learn about their own Behavior*. Report HPP-84-16, Department of Computer Science, Stanford University, 1984.
9. Hayes-Roth, B.: *A Blackboard Architecture for Control*. Artificial Intelligence 26:251-321, 1985.
10. Hayes-Roth, B. and Hewett, M.: *Learning Control Heuristics in BB1*. Report HPP-85-2, Department of Computer Science, 1985.
11. Hayes-Roth, B., Buchanan, B.G., Lichtarge, O., Hewett, M., Altman, R., Brinkley, J., Cornelius, C., Duncan, B., and Jardetzky, O.: *PROTEAN: Deriving protein structure from constraints*. Proceedings of the AAAI, 1986, p. 904-909.
12. Jardetzky, O.: *A Method for the Definition of the Solution Structure of Proteins from NMR and Other Physical Measurements: The LAC-Repressor Headpiece*. Proceedings of the International Conference on the Frontiers of Biochemistry and Molecular Biology, Alma Alta, June 17-24, 1984, October, 1984.
13. Lichtarge, Olivier: *Structure determination of proteins in solution by NMR*. Ph.D. Thesis, Stanford University, November, 1986.
14. Lichtarge, Olivier, Cornelius, Craig W., Buchanan, Bruce G., Jardetzky, Oleg: *Validation of the First Step of the Heuristic Refinement Method for the Derivation of Solution Structures of Proteins from NMR Data*, April 1987. Submitted to Proteins: Structure, Function, and Genetics.

National AIM Project: KNOWLEDGE ENGINEERING FOR
 RADIATION THERAPY

Principal Investigator: Ira J. Kalet, Ph.D.
 School of Medicine
 University of Washington at Seattle
 Seattle, Washington 98195
 (206) 548-4107

We are developing an expert system for planning of radiation therapy for head and neck cancers. The project will ultimately combine knowledge-based planning with numerical simulation of the radiation treatments. The numerical simulation is needed in order to determine if the proposed treatment will conform to the goals of the plan (required tumor dose, limiting dose to critical organs). The space of possible radiation treatments is numerically very large, making traditional search techniques impractical. Yet, with modern radiation therapy equipment, the design of treatment plans might be significantly aided by automatically generating plans that meet the treatment constraints. The project will result in systematization of knowledge about radiation treatment design, and will also provide an example of how to represent and solve design problems with a knowledge based system.

This project has some relevance to computer science as well, in that our approach, if successful, may contribute to a better understanding of design problem solving with knowledge-based systems.

REFERENCES

1. I. Kalet and W. Paluszynski: A Production Expert System for Radiation Therapy Planning. Proceedings of the AAMSI Congress 1985, May 20-22, 1985, San Francisco, California. Edited by Allan H. Levy and Ben T. Williams. American Association for Medical Systems and Informatics, Washington, D.C., 1985.
2. W. Paluszynski and I. Kalet: Radiation Therapy Planning: A Design Oriented Expert System. WESTEX-87 (Western Conference on Expert Systems), Anaheim, California, June 2-4, 1987.
3. I. Kalet and J. Jacky: Knowledge-based Computer Simulation for Radiation Therapy Planning. Proceedings of the Ninth International Conference on the use of Computers in Radiotherapy, Scheveningen, the Netherlands, June 1987. North Holland, 1987.
4. I. Kalet and W. Paluszynski: Knowledge-based Computer Systems for Radiotherapy Planning. American Journal of Clinical Oncology. In press.

Stanford Project: RADIX -- DERIVING KNOWLEDGE FROM
TIME-ORIENTED CLINICAL DATABASES

PENGUIN -- Applying Database and Knowledge base
Technology to Medical Instrumentation

Principal Investigators: Gio C.M. Wiederhold, Ph.D.
Departments of Computer Science
and Medicine
Stanford University
Stanford, California 94305
(415) 497-0685 (WIEDERHOLD@SUMEX-AIM)

Thierry Barsalou, M.D.
Medical Computer Science
Stanford University
(BARSALOU@SUMEX-AIM)

Robert L. Blum, M.D., Ph.D.
Stanford University
(BLUM@SUMEX-AIM)

The RADIX research has been phased out during this year. RADIX had two main goals: (1) to explore the usefulness of knowledge-based techniques to derive medical knowledge from clinical database (DB) systems containing non-randomized, non-protocol patient observations and (2) to develop a program and set of techniques for automated summarization of patient records. The process of reliably deriving causal relationships has proven to be quite difficult because of the complexity of disease states and time relationships, strong sources of bias, and problems of missing and outlying data. However, the experience gained in the RADIX project and its predecessor project, RX, continues to influence basic research directions. The problem of automated knowledge acquisition remains an important area of research, and we expect that the foundations laid here will influence work of others as well.

PENGUIN is a project aimed at developing methods for bringing together expert system and database technologies in an integrated advice system. Databases and expert systems share a common goal -- generating useful information for action -- but accomplish their tasks separately, using different principles. It is clear, however, that future information systems will require both the problem-solving capabilities of expert systems (ESs) and the data-handling capabilities of database management systems (DBMSs). Indeed, combining database and expert system technologies into expert database systems (EDSs) is an emerging research area. One can define an EDS as "a system for developing applications requiring knowledge-directed processing of shared information". From a perspective of developing advanced biomedical information systems, this definition conveys two goals: (1) enhancing DBMSs with structuring and manipulation tools that take more semantics into account; (2) allowing ESs to access and to handle efficiently information stored in database(s). In this research project, we investigate the hypothesis that the object-oriented approach to knowledge and information representation can serve as a unifying scheme for developing EDSs.

We explore this hypothesis in a practical biomedical environment, Fluorescence Activated Cell Sorting (FACS). FACS is emerging as a major source of information for biomedical research and clinical practice. Currently, achieving good FACS performance requires analyzing and integrating various complex data and knowledge sources in order to fulfill the information needs of FACS investigators. This work

involves a close collaboration between researchers in the Medical Information Sciences Program and the Departments of Computer Science and Genetics.

REFERENCES

PENGUIN PUBLICATIONS

1. Barsalou, Thierry and Gio Wiederhold: "Automating a Cell Counter"; to be published in The International Journal of Artificial Intelligence in Engineering, Computational Mechanics Publ., UK, 1988.
2. Barsalou, Thierry and Gio Wiederhold: "Applying a Semantic Model to an Immunology Database"; in W.W. Stead (editor), Proceedings of the Eleventh Symposium on Computer Applications in Medical Care, pages 871-877, IEEE Computer Society Press, Washington, D.C., November 1987.
3. Barsalou, Thierry, W.A. Moore, L.A. Herzenberg, and G. Wiederhold: "A Database System to Facilitate the Design of FACS Experiment Protocols" (abstract); Cytometry, Vol.97, August 1987.
4. Barsalou, T. and G. Wiederhold. *Knowledge-based mapping of relations into objects*. To appear in the International Journal of AI in Engineering, 1988.
5. Barsalou, T. *An object-based architecture for biomedical expert database systems*. Submitted to the IEEE Twelfth Symposium on Computer Applications in Medical Care, 1988.

RADIX PUBLICATIONS

Monograph

1. Blum, R.L.: *Discovery and representation of causal relationships from a large time-oriented clinical database: The RX project*. in D.A.B. Lindberg and P.L. Reichertz (Eds.), LECTURE NOTES IN MEDICAL INFORMATICS, Vol. 19, Springer-Verlag, New York, 1982.

Journal Articles

1. Blum, R.L.: *Computer-Assisted Design of Studies using Routine Clinical Data: Analyzing the Association of Prednisone and Cholesterol*. Annals of Internal Medicine 104(6):858-868, June, 1986.
2. Blum, R.L.: *Discovery, confirmation, and incorporation of causal relationships from a large time-oriented clinical database: The RX Project*. Computers and Biomedical Research 15(2):164-187, April, 1982.
3. Walker, M.G. *How Feasible is Automated Discovery?* IEEE Expert 2(1):70-82, Spring, 1987.

Conference Proceedings

1. DeZegher-Geets, I.M., Freeman, A.G., Walker, M.G., Blum, R.L., and Wiederhold, G.C.M. *Computer-aided Summarization of a Time-oriented Medical Data Base*. In Proceedings of the Third Annual Conference on Computerization of Medical Records. Institute for Medical Record Economics, 1987.

2. Downs, S., Walker, M.G., and Blum, R.L. *Automated Summarization of On-line Medical Records*. In Proceedings of the Fifth World Congress on Medical Informatics (Medinfo), pages 800-804. Elsevier Science Publishers, 1986.
3. Walker, M.G., and Blum, R.L. *Towards Automated Discovery from Clinical Databases: the RADIX Project*. In Proceedings of the Fifth World Congress on Medical Informatics (Medinfo), pages 32-36. Elsevier Science Publishers, 1986.

Stanford Project: REFERENCE Project

Principal Investigators: Bruce G. Buchanan, Principal Investigator
Computer Science Department
Stanford University
Stanford, California 94305

Byron W. Brown, Co-Principal Investigator
Department of Medicine
Stanford University Medical Center

Daniel E. Feldman, Associate Investigator
Department of Medicine
Stanford University Medical Center

The goals of this project are related both to medical science and artificial intelligence: (a) use AI methods to allow the informed but non-expert reader of the medical literature to evaluate a randomized clinical trial, and (b) use the interpretation of the medical literature as a test problem for studies of knowledge acquisition and fusion of information from disparate sources. REFERENCE and REVIEWER, a planned extension, will be used to evaluate the medical literature of clinical trials to determine the quality of a clinical trial, make judgments on the efficacy of the treatment proposed, and synthesize rules of clinical practice. The research is an initial step toward a more general goal - building computer systems to help the clinician and medical scientist read the medical literature more critically and more rapidly for use in making clinical decisions.

REFERENCES

1. Haggerty, J.: *REFERENCE and RULECRITIC: Two prototypes for assessing the quality of a medical paper*. REPORT KSL-84-49. Master's Thesis, Stanford University, May 1984.
2. *Chavez, R. Martin and Cooper, G. F.: *KNET: Integration Hypermedia and Normative Bayesian Modeling*. REPORT KSL. Stanford University, March 1988.
3. *Lehmann, H. *Knowledge Acquisition for Probabilistic Expert Systems*. Submitted to Symposium on Computer Applications in Medical Care, 1988.

References

1. Coulter, C. L. "Research Instrument Sharing." *Science* 201, 4354 (1978).
2. Lederberg, J. "Digital Communications and the Conduct of Science: The New Literacy." *Proceedings of the IEEE* 66, 11 (1978).
3. NLM Planning Panel 4. NLM Long Range Plan: Medical Informatics. Shortliffe, E.H., Panel 4 Chairman , National Library of Medicine, January, 1987.
4. Scheifler, R.W. and Gettys, J. The X Window System. Draft October 1986. To appear in a special issue of the *ACM Transactions on Graphics* -- user interface software
5. Steele, Guy L., Jr.. *Common Lisp - The Language*. Digital Press, Burlington, MA, 1984.